

---

## EXPERT SYSTEMS AND THE ENVIRONMENT

---

### History of expert systems

Expert systems are computer programs that can perform some task which typically requires the capabilities of a skilled human. These tasks are usually of a decision-making nature rather than physical actions. Examples of such tasks are managing water levels in a wetland, forecasting weather conditions, assessing environmental impacts, and selecting mitigation measures for environmental hazards. As computer programs that contain human expertise, they are referred to variously by the labels expert systems, knowledge-based systems, inference systems or rule-based systems.

Expert systems have evolved as a highly commercializable offshoot of research in the subfield of computer science called *artificial intelligence* (AI). Since its unofficial inception at the Dartmouth Summer Research Project on Artificial Intelligence in 1956 (attended by illuminaries such as Marvin Minsky, Allen Newell, Herbert Simon, Claude Shannon and John McCarthy), AI has had as one of its primary goals the creation of 'thinking machines.' While this ambitious goal has not yet been attained to anyone's acknowledgment, there have been substantial advances in what we now know about human thinking and learning. Along the way, research in AI from the late 1950s to the 1970s at Stanford, MIT and Carnegie-Mellon Universities provided some very powerful techniques for codifying human experience and knowledge so that computers can store it and apply it to solve practical problems. The mid-1970s saw the emergence of the first expert systems for applications such as medical diagnosis (*Mycin*, by Shortliffe), chemical data analysis (*Dendral*, by Lindsay and others), and mineral exploration (*Prospector*, by Duda and others). For further information, see Barr and Feigenbaum (1982). Since that time, the

proliferation of this technology and the need to extend human expertise beyond the local time and place of the expert have led to the development of thousands of expert systems across hundreds of different fields.

### How expert systems work

Typically a user interacts with an expert system in consultation dialog (Figure E29), much like one would converse with a human expert. The user explains the problem to be solved, provides necessary background information and queries the system about proposed solutions. In the knowledge acquisition mode, a human expert interacts with the system to create a knowledge base of what he or she knows in a particular subject area. Through these two operational modes the expert system acts, in some sense, like an intermediary between the expert (acquisition mode) and the user (consultation mode).

Most expert systems consist of several distinct components. These are knowledge base, working memory, reasoning engine, explanation subsystem and a user interface. The *knowledge base* contains the scientific knowledge and experience for the particular area of expertise. Imagine that we are designing an expert system to diagnose automobile engine malfunctions. We might want to include knowledge about spark plugs, fuel pump, battery, starter, fuel injectors, etc., and also how these engine components affect engine operation. A competent mechanic can usually pinpoint engine problems fairly quickly with only a small amount of information about the functioning of the various parts. Often a specialist, such as a mechanic, possesses intuition that he or she has acquired through years of experience. This intuition is often reified in rules-of-thumb (or good guesses) that allow the specialist to solve problems quickly and effectively. For this type of expert knowledge to be used by a computer it must be represented in some way that the computer can easily manipulate. There are numerous techniques for *knowledge representation*, but traditionally the most common one is the use of condition-action rules (see

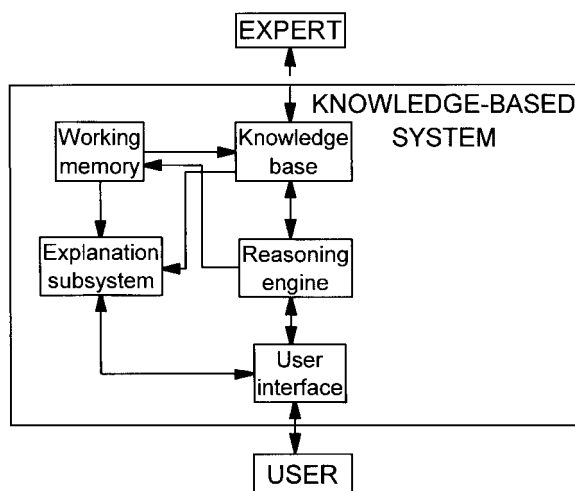
Luger and Stubblefield, 1989, for a comprehensive review of these techniques). Condition-action rules are IF-THEN statements where the consequent action(s) are performed if the premise conditions are true. For example, IF battery charged AND battery-cables = clean AND engine-starting = not cranking THEN check starter. This method of knowledge representation is popular because each rule is modular and contains a 'chunk of domain knowledge, expert system programmers find rules easy to program, and experts are often able to express their heuristic knowledge in the IF-THEN format.

*Working memory* is like the short-term memory of the expert system. It contains assertions about the problem currently under investigation. These assertions may be obtained from the user (via queries), from external programs, from a real-time process, or from external data files. Assertions may be facts gathered from the above sources, or they may be hypotheses which have been inferred from other facts that are already known. Because the ultimate goal of knowledge system consultation is to infer problem solutions, some of these intermediate hypotheses will eventually be solutions. All facts and hypotheses in the working memory together describe the current context, or the current state, of a consultation session. Usually a closed world assumption is assumed, i.e., only those assertions that are present in the working memory are true and all other possible assertions about the state of the world are assumed false.

While the knowledge base and working memory are passive entities, the *reasoning engine* navigates through the knowledge base and registers established assertions in the working memory. A reasoning engine operating on a knowledge base and working memory is how an expert system solves problems. Navigation is performed by the particular control strategy that the reasoning engine employs. A control strategy determines the order in which knowledge base elements (such as rules) are examined in order to arrive at the solution to a problem. Assertions are established as true by the particular inferencing mechanism used. In a rule-based knowledge representation, the inferencing method is usually *modus ponens* and rules are selected for evaluation either by the content of their premise conditions (data-driven control) or by their consequent actions (goal-driven control). Details of how the reasoning engine operates are determined by the knowledge representation method used, what types of assertions must be made, and the overall problem-solving methods that are applied.

The purpose of an *explanation subsystem* is to enable the expert system to display to users an understandable account of the motivation for all of its actions and conclusions. Explanation is part of the larger issue of human factors engineering, which also includes the user interface – i.e., the hows and whys of a computer system's interaction with users. Explanation systems are not involved with the correct execution of an expert system. Instead, their purpose is to convince the user that the system's conclusions are reasonable, to explain how it reached those conclusions, and to aid system developers in debugging the knowledge base and the reasoning methods.

The term *user interface* refers to the physical and sensory interaction between computer and user. Functionally, this means how the user inputs information to the system and how information is returned to the user. The more natural (i.e., intuitive and understandable) this interface is, the more effective the humancomputer interaction will be. Traditionally, this interaction has been serial and text based using the conventional, interactive terminal format. Recent advances in computer interfaces enable expert systems to utilize display



**Figure E29** The expert system operates either in consultation mode or knowledge acquisition mode. The various system components enable it to solve problems for which it has knowledge in the knowledge base, to interact with users, and to explain the rationale for the solutions it reaches. Information flows are depicted with arrows.

graphics, hot graphics (graphical objects that perform some action when activated), point-and-click operations, video, sound and animation. For most software users, the interface is the application, and hence expert systems may fall into disuse if they lack good user-interface capabilities.

### Environmental science applications

Limitations of space prohibit enumeration of all the expert systems developed in environmental science as it is broadly defined. An alternative way to present them is functionally, i.e., according to the types of problems that they address. The non-exclusive categories that seem to capture most applications are classification, prediction, interpretation, planning, monitoring and control, and analysis. The categorical approach is advantageous because the reader then acquires an appreciation of the broad applicability of expert system methodology without becoming distracted by details that are specific to particular applications (see the bibliography in Davis and Clark, 1989, and the surveys in Hushon, 1987, and Moninger and Dyer, 1988).

*Classification* problems are the most common type of application. This is due in part to our inherent human need to classify objects and events as being members of particular groupings. A salient characteristic of classification problems is that there is a finite (usually small) and enumerable list of possible groups; this makes these problems relatively easy to solve. Hence, all problems that fall into a particular solution group are treated similarly with respect to action. Diagnosis is a very common application problem, where systems are diagnosed in terms of the causes of malfunction. These include biological systems (e.g., trees, crops or fish populations), hydrological and chemical systems (e.g., lakes and streams), mechanical systems (e.g., waste treatment) or physical systems (e.g., hailstorm severity). The cause may be a pathogen, a malfunctioning pump, a parasite, a climate change, and so on. Other non-diagnostic classification systems only seek to place an object or event into a particular category without labeling that category as malfunctioning; for example, identification of type of atmospheric inversion, classification of soils, selection of options in silviculture or of insecticides, or identification of species.

Another large class of expert systems applications includes those that deal with *prediction*. These estimate some important future characteristic of an environmental system based on current details about it. Some examples of prediction problems are forecasting for weather and other environmental phenomena, qualitative modeling of biological or physical systems (e.g., vegetation change, crop production and wildlife populations), and damage estimation (e.g., following toxic contamination, for insect epidemics or for flooding). When these expert systems select their predictions from a small set of possible future conditions, they can also be categorized as classification expert systems. It should be apparent that there is some overlap between classification and prediction problems. In fact, all these categories are non-exclusive, and hence overlaps exist between most of them. In fact, many systems can be categorized in multiple ways.

*Interpretation* problems are similar to prediction problems except that the characteristic to be estimated is a current one, rather than a future one. Because this characteristic condenses and summarizes the information about an environmental system, it usually carries with it some important management

implications. Ways in which expert systems have been applied include hazard and risk ratings (e.g., fire danger rating, and contamination or toxicity potential estimation), environmental assessment (e.g., impacts of human intervention, cost estimation, and report evaluation or generation), data interpretation (e.g., model interpretation, site selection or ranking, species selection and equipment selection), and management actions (e.g., fire suppression, and crop production and treatment prescriptions).

Solutions to the above three categories of problems most often consist of a single action or parameter estimate. *Planning* type problems, on the other hand, are resolved by specifying an ordered set of actions to be performed. Because a large number of possible action sequences is possible, planning problems tend to be much more difficult to solve and are more computationally costly. Examples of reported applications in this area are catastrophe mitigation (e.g., hazardous site cleanup, and fire suppression), forest and agriculture production (planting, treatment and harvest), construction (e.g., roads or airport runways), and scheduling and resource planning (e.g., for regional water quality, landscape and land use). Expert systems provide a viable approach to solving planning problems because these problems usually have a fairly well-defined goal that is constrained by certain of their attributes. Moreover, they are non-quantitative in nature and require a systematic search through a large number of possible solutions.

In contrast to the off-line decision making that is inherent in the problems described above, there are situations in which decisions are made as part of real-time operations. *Monitoring and control problems* are of this type. In many of these instances monitoring and control activities are intertwined in the sense that a process is monitored by an expert system that also takes action when some condition signals its attention. At other times, an expert system only performs monitoring, and a human being performs the control action. Examples of monitoring and control applications are very few in the environmental sciences, and this category is only mentioned here for the sake of completeness.

A final application for expert systems in environmental science is in the area of *analysis*. Here, an expert system assists with evaluation of a system, or data about a system, or it enhances the operation of existing analysis methods. In the first case, expert systems can help collect or filter data, or suggest analyses for data; in the latter case they serve as 'intelligent' front ends or internal enhancements to existing software. Expert systems appear as laboratory recording aides, report generators, data collection and selection aides, cartographic aides, data error detectors and correctors, curve shape analyzers, and data quality assessors. As intelligent front ends and imbedded 'intelligence,' expert systems have been used with ecological models, geographic information systems, remote sensing and cartographic systems. Most of these systems are designed for in-house laboratory use to enable scientists and technicians to work better and more efficiently.

The actual deployment of environmental science expert systems has been meager. Most expert systems have been developed at universities and other research laboratories. Consequently, there is often little incentive for developers to translate their work from laboratories to operational settings. Also, for some scientific disciplines the user group for these expert systems – i.e., the field personnel or practitioners – has not completely adopted information technology. Still other questions remain about software maintenance and technical

support, which are time-consuming tasks that developers are often unwilling or unable to assume. No survey has been done to estimate the ratio of delivered to developed systems, but for the several hundred systems that have been developed to a reasonable degree of completeness, it is probably accurate to say that no more than a few dozen are actually used on a regular basis.

Daniel L. Schmoldt

### Bibliography

- Barr, A., and Feigenbaum, F.A., 1982. *The Handbook of Artificial Intelligence*, Volume 2. Los Altos, Calif.: William Kaufmann, 428 pp.
- Davis, J.R., and Clark, J.L., 1989. A selective bibliography of expert systems in natural resource management. *AI Applications*, **3**, 1-18.
- Hushon, J.M., 1987. Expert systems for environmental problems. *Environ. Sci. Technol.*, **21**, 838-41.
- Luger, G.F., and Stubblefield, W.D., 1989. *Artificial Intelligence and the Design of Expert Systems*. Redwood City, Calif.: Benjamin Cummings.
- Moninger, W.R., and Dyer, R.M., 1988. Survey of past and current AI work in the environmental sciences. *AI Applications*, **2**, 48-52.

### Cross-references

Environmental Audit  
Environmental Impact Analysis (EIA), Statement (EIS)  
Environmental Statistics  
Systems Analysis

ENCYCLOPEDIA OF EARTH SCIENCES SERIES

# **ENCYCLOPEDIA *of* ENVIRONMENTAL SCIENCE**

*edited by*

**DAVID E. ALEXANDER**

University of Massachusetts

*and*

**RHODES W. FAIRBRIDGE**

NASA – Goddard Institute for Space Studies



**KLUWER ACADEMIC PUBLISHERS**  
DORDRECHT | BOSTON | LONDON

A C.I.P. Catalogue record for this book is available from the Library of Congress.

ISBN 0-412-74050-8

---

Published by Kluwer Academic Publishers  
PO Box 17, 3300 AA Dordrecht, The Netherlands

Sold and distributed in North, Central and South America  
by Kluwer Academic Publishers, PO Box 358,  
Accord Station, Hingham, MA 02018-0358, USA

In all other countries, sold and distributed  
by Kluwer Academic Publishers  
PO Box 322, 3300 AH Dordrecht, The Netherlands

Every effort has been made to contact the copyright holders of the figures and tables which have been reproduced from other sources. Anyone who has not been properly credited is requested to contact the publishers, so that due acknowledgement may be made in subsequent editions,

All Rights Reserved

© 1999 Kluwer Academic Publishers

No part of this publication may be reproduced or utilized in any form or by any means, electronic, mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission from the copyright owner.

Printed and bound in Great Britain by MPG Books, Bodmin, Cornwall.